

Technical Note: Convergence analysis of a polyenergetic SART algorithm

T. Humphries*

Department of Mathematics, Oregon State University

(Dated: May 15, 2015)

Purpose: We analyze a recently proposed polyenergetic version of the simultaneous algebraic reconstruction technique (SART). This algorithm, denoted pSART, replaces the monoenergetic forward projection operation used by SART with a post-log, polyenergetic forward projection, while leaving the rest of the algorithm unchanged. While the proposed algorithm provides good results empirically, convergence of the algorithm was not established mathematically in the original paper.

Methods: We analyze pSART as a nonlinear fixed point iteration by explicitly computing the Jacobian of the iteration. A necessary condition for convergence is that the spectral radius of the Jacobian, evaluated at the fixed point, is less than one. A short proof of convergence for SART is also provided as a basis for comparison.

Results: We show that the pSART algorithm is not guaranteed to converge, in general. The Jacobian of the iteration depends on several factors, including the system matrix and how one models the energy dependence of the linear attenuation coefficient. We provide a simple numerical example that shows that the spectral radius of the Jacobian matrix is not guaranteed to be less than one. A second set of numerical experiments using realistic CT system matrices, however, indicates that conditions for convergence are likely to be satisfied in practice.

Conclusion: Although pSART is not mathematically guaranteed to converge, our numerical experiments indicate that it will tend to converge at roughly the same rate as SART for system matrices of the type encountered in CT imaging. Thus we conclude that the algorithm is still a useful method for reconstruction of polyenergetic CT data.

I. INTRODUCTION

The algebraic reconstruction technique (ART) or Kaczmarz' method [1, 2] is a well-known method for approximately solving systems of linear equations,

$$A\mathbf{x} = \mathbf{b}, \quad (1)$$

where \mathbf{x} is a column vector of size n , \mathbf{b} is a column vector of size m , and A is $m \times n$. The method has a long history in CT image reconstruction, in which \mathbf{b} represents the post-log measured projection data, \mathbf{x} represents the image to be reconstructed, and the (i, j) th entry of A represents the length or area of intersection of the i th ray with the j th pixel of the image. The values in \mathbf{x} represent the averaged linear attenuation coefficient (LAC) within each pixel, and are usually constrained to be positive. Since the model is linear, it is implicitly assumed that the X-ray beam used to generate the data is monoenergetic, and the reconstructed values correspond to the LAC of the tissue at that energy.

Beginning from an initial estimate $\mathbf{x}^{(0)}$, ART iteratively projects the current iterate, $\mathbf{x}^{(k)}$, onto the hyperplane defined by one of the m equations defined by (1). The ART iteration is given by

$$\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} - \frac{\langle a_i, \mathbf{x}^{(k)} \rangle - b_i}{\langle a_i, a_i \rangle} a_i^T, \quad (2)$$

where a_i is the i th row of A , $\langle \cdot, \cdot \rangle$ denotes the dot product, and i is chosen to be $(k \bmod m) + 1$ in the classical

version of the algorithm. Some references consider this step to be a sub-iteration, and a full iteration of ART to consist of successively applying this iteration for all m equations. If the system is consistent, the iteration is guaranteed to converge to a solution, while in the inconsistent case the sub-iterations converge to a limit cycle [3].

The simultaneous ART (SART) method [4, 5] is a variant of ART in which the corrections generated by the ART sub-iterations (2) are combined and applied simultaneously. The iteration can be expressed concisely in terms of matrix operations as

$$\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} - DA^T M (A\mathbf{x}^{(k)} - \mathbf{b}), \quad (3)$$

where D and M are diagonal matrices, with

$$D_{jj} = \frac{1}{\beta_j}, \quad \beta_j = \sum_{k=1}^m |a_{kj}|, \quad j = 1 \dots n$$

$$M_{ii} = \frac{1}{\gamma_i}, \quad \gamma_i = \sum_{k=1}^n |a_{ik}|, \quad i = 1 \dots m. \quad (4)$$

In other words, β_j is the 1-norm of the j th column of A , and γ_i is the 1-norm of the i th row.

The primary advantage of SART over ART is that SART is less sensitive to noisy data [4]. It should be noted that the iteration (3) computes an update using projection data corresponding to all views simultaneously, while the original SART algorithm [4] computes a sequence of updates using projection data corresponding to only one view of the object at one time, to accelerate convergence. For simplicity of analysis we will only consider equation (3), as this was the used as the basis for

* Email: humphrit@science.oregonstate.edu

the polyenergetic algorithm that we will analyze. Like ART, SART has been proven to converge [6, 7]; see also the Appendix to this paper.

As mentioned, both ART and SART implicitly assume that the data are generated from a monoenergetic X-ray beam, i.e. that

$$p_i = I \exp(-\langle a_i, \mathbf{x} \rangle), \quad (5)$$

where p_i is the measured intensity of the i th beam, and I is the blank scan intensity (assumed to be independent of i). Taking the log of the data and rearranging terms then gives a linear system equivalent to (1).

$$-\ln \frac{p_i}{I} = \langle a_i, \mathbf{x} \rangle. \quad (6)$$

In practice, however, the X-rays generated by clinical CT hardware are usually polyenergetic. A typical model for polyenergetic X-ray measurements is

$$p_i = \int I(\varepsilon) \exp(-\langle a_i, \boldsymbol{\mu}(\varepsilon) \rangle) d\varepsilon, \quad (7)$$

where ε refers to the energy of an incident x-ray, $I(\varepsilon)$ is the initial intensity of the beam corresponding to that energy (i.e. the beam's spectrum) and $\boldsymbol{\mu}(\varepsilon)$ is a vector of attenuation coefficients, whose values depend on X-ray energy. This system of equations can no longer be linearized, and it is well-known that reconstructing an image from this data using conventional means (such as ART, SART, or filtered back projection) produces an image containing beam hardening artifacts [8]. This motivates the need for polyenergetic iterative reconstruction algorithms, e.g. [9–18].

The polyenergetic SART (pSART) algorithm [18] was recently proposed as one such reconstruction technique. In this approach, the beam spectrum is discretized into h energy levels ε_h , with weighting terms I_h computed to approximate the continuous spectrum. The vector-valued function $\boldsymbol{\mu}(\varepsilon)$ is then modeled as a function of a vector \mathbf{t} , representing the attenuation map of the object at a reference energy, ε_0 , which was chosen to be 70 keV. This function makes use of tabulated energy-dependent LAC values for some suitable reference materials, such as air, fat, breast, soft tissue and bone. Letting t_j denote the LAC value for pixel j at the reference energy, the LAC of that pixel for all other energies ε is then given by

$$\mu(t_j, \varepsilon) = \frac{[\mu_{k+1}(\varepsilon_0) - t_j]\mu_k(\varepsilon) + [t_j - \mu_k(\varepsilon_0)]\mu_{k+1}(\varepsilon)}{\mu_{k+1}(\varepsilon_0) - \mu_k(\varepsilon_0)}, \quad (8)$$

where $\mu_k(\varepsilon)$ and $\mu_{k+1}(\varepsilon)$ are the tabulated, energy-dependent LAC functions for the two base materials with LAC values adjacent to t_j at the reference energy. So for instance, if the value of t_j is between the LAC for soft tissue and the LAC for bone at the reference energy, then its LAC at all other energies is obtained by linear interpolation between the corresponding values for bone and

soft tissue, with the weighting determined by the values at the reference energy.

One can then define a polyenergetic forward projection operator, $\mathcal{P} : \mathbb{R}^n \rightarrow \mathbb{R}^m$, which acts on \mathbf{t} :

$$[\mathcal{P}(\mathbf{t})]_i = \sum_h I_h \exp(-\langle a_i, \boldsymbol{\mu}(\mathbf{t}, \varepsilon_h) \rangle). \quad (9)$$

The pSART iteration is then defined as

$$\mathbf{t}^{(k+1)} = \mathbf{t}^{(k)} - DA^T M \left(-\ln [\mathcal{P}(\mathbf{t}^{(k)})] + \ln(\mathbf{p}) \right), \quad (10)$$

with D and M defined as in (4). Equations (9) and (10) are equivalent to equations (14) and (15) in Ref. 18, although our notation is somewhat different. The only difference from the SART iteration (3) is that the log of the monoenergetic forward projection has been replaced by the log of the polyenergetic forward projection. The algorithm produces a single attenuation map of the object, \mathbf{t} , with LAC values corresponding to the reference energy.

In Ref. 18 the authors state that this modification solves the problem of inconsistency between the polyenergetic data and the monoenergetic model implicitly assumed by the conventional SART approach. While it is clear that a vector \mathbf{t} satisfying $\mathbf{p} = \mathcal{P}(\mathbf{t})$ is a fixed point of this iteration, this does not guarantee convergence of the algorithm. Experimental results indicated that the method is effective, however. In the next section we analyze the convergence of pSART as a fixed point iteration.

II. CONVERGENCE OF PSART

We first consider the SART iteration (3). The iteration can be written in the form

$$\mathbf{x}^{(k+1)} = T\mathbf{x}^{(k)} + \mathbf{c}, \quad (11)$$

where $T = I - DA^T M A$ and $\mathbf{c} = DA^T M \mathbf{b}$. One can show that the spectral radius of T (the magnitude of its largest eigenvalue), denoted $\rho(T)$, is strictly less than 1 (see Appendix). This guarantees convergence of the algorithm to a solution, if the system is consistent. It has been proven that if no exact solution exists, SART converges to a weighted least-squares solution [6, 7].

The pSART algorithm is a nonlinear fixed point iteration. We write the iteration as

$$\mathbf{t}^{(k+1)} = \mathbf{t}^{(k)} - DA^T M f(\mathbf{t}^{(k)}), \quad (12)$$

$$\equiv F(\mathbf{t}^{(k)}) \quad (13)$$

where $f(\mathbf{t}) = -\ln [\mathcal{P}(\mathbf{t})] + \ln(\mathbf{p})$. Assuming that a solution to the system of nonlinear equations exists (which corresponds to a fixed point of the iteration), a necessary condition for convergence is that the spectral radius of the Jacobian matrix of F , $\rho(J_F)$, must be less than one when evaluated at the solution. Note that this does not

guarantee that the algorithm converges to the solution from any starting point; simply that it cannot converge if the condition does not hold. Let \mathbf{t}^* denote a solution to the nonlinear system of equations $\mathbf{p}_i = [\mathcal{P}(\mathbf{t})]_i$. A straightforward calculation gives

$$J_F(\mathbf{t}^*) = I - DA^T M J_f(\mathbf{t}^*), \quad (14)$$

where the (i, j) th element of the $m \times n$ Jacobian matrix $J_f(\mathbf{t})$ is given by

$$\frac{\partial f_i}{\partial t_j}(\mathbf{t}) = \frac{a_{ij}}{[\mathcal{P}(\mathbf{t})]_i} \left[\sum_h I_h \exp(-\langle a_i, \boldsymbol{\mu}(\mathbf{t}, \varepsilon_h) \rangle) \frac{\partial \mu}{\partial t}(t_j, \varepsilon_h) \right]. \quad (15)$$

It follows from (8) that

$$\frac{\partial \mu}{\partial t}(t, \varepsilon) = \frac{\mu_{k+1}(\varepsilon) - \mu_k(\varepsilon)}{\mu_{k+1}(\varepsilon_0) - \mu_k(\varepsilon_0)}, \quad (16)$$

where μ_{k+1} and μ_k again refer to the base material LACs that are adjacent to t at the reference energy ε_0 . This implies that for a fixed value of ε , this partial derivative is piecewise constant, with discontinuities when t is equal to one of the reference LACs at ε_0 .

A general analysis of the spectral radius of $J_F(\mathbf{t}^*)$ is difficult as it has a complicated dependence on the system matrix A , the spectrum I_h , and the choice of base materials. We now show with a numerical experiment that one cannot guarantee that $\rho(J_F(\mathbf{t}^*)) < 1$, in general.

We consider a simple example where $m = n = 2$, illustrated in Fig. 1. The object consists of two pixels of size 1×1 cm, with LACs of $t_1 = 0.1 \text{ cm}^{-1}$ and $t_2 = 0.16 \text{ cm}^{-1}$ at a reference energy of 70 keV. We first consider the case of a monoenergetic 70 keV beam with intensity $I = 1$. The first beam travels through both pixels horizontally, while the second beam has a length of intersection of roughly 0.28 cm with the first pixel and 1.13 cm with the second pixel. After taking logarithms, we obtain the 2×2 linear system of equations

$$\begin{bmatrix} 1 & 1 \\ 0.28 & 1.13 \end{bmatrix} \begin{bmatrix} t_1 \\ t_2 \end{bmatrix} = \begin{bmatrix} 0.260 \\ 0.209 \end{bmatrix}, \quad (17)$$

which can be solved using ART or SART. Fig. 2 shows the progression of both ART and SART for this system, starting from an initial guess of zero. One can see that ART sequentially projects onto the two lines defining the equations, while the iterates generated by SART follow a path between the two equations.

We now consider the case of a polyenergetic x-ray spectrum. The LACs of the two pixels at the reference energy determine the LACs at all other energies ε according to (8). The discrete spectrum used in this experiment consists of eleven energy bins, obtained from a continuous 130 kVp spectrum generated using the Siemens Spektrum online tool [19, 20]. This spectrum and the attenuation curves for the two materials are shown in Fig. 3. The attenuation values for the reference materials were




FIG. 1. Example problem consisting of two pixels.




FIG. 2. Convergence of ART & SART for the monoenergetic experiment. The black lines correspond to the system of equations (17), red line correspond to the sequence of iterates generated by ART, blue line to the sequence of iterates generated by SART.

obtained from Ref. 21. Under the polyenergetic model, we obtain a nonlinear system of equations:

$$\begin{aligned} p_1 &= \sum_h I_h \exp(-\mu(t_1, \varepsilon_h) - \mu(t_2, \varepsilon_h)) \\ &\approx \exp(-0.314) \\ p_2 &= \sum_h I_h \exp(-0.28\mu(t_1, \varepsilon_h) - 1.13\mu(t_2, \varepsilon_h)) \\ &\approx \exp(-0.253) \end{aligned}$$

The two beams undergo more attenuation than in the monoenergetic experiment (17) because the spectrum contains a higher proportion of X-rays with energies less than 70 keV. In Fig. 4, this nonlinear system of equations is illustrated with black curves. As either t_1 or t_2 increases, the LAC value of the pixel at lower energies increases rapidly, meaning that the coefficient in the other pixel must decrease rapidly to compensate. Thus the curves are bent. For the sake of illustration, we have implemented a polyenergetic version of ART (denoted pART) that is analogous to pSART. The red and blue lines show the progression of the pART and pSART iterations, respectively. For this experiment (left figure), one can see that both iterations converge to the solution of the nonlinear system of equations.

We now give a case where the iteration fails to converge. In the right figure of Fig. 4, the progress of the two iterations is shown for a slightly modified experiment, where t_2 was changed from 0.16 cm^{-1} to 0.24 cm^{-1} . All




FIG. 3. X-ray spectrum and attenuation curves used for the polyenergetic experiment. Left: Continuous spectrum (blue line) and discrete energies (red crosses) used for the summation. The spectrum has been normalized to have an integral of 1. Right: Attenuation curves for the base materials as well as the interpolated curves for values of $t = 0.1, 0.16$ and 0.24 (red dashed lines). The reference energy of 70 keV is indicated by the dashed black line.




FIG. 4. Convergence of pART and pSART for two polyenergetic experiments. The black curves correspond to the nonlinear system of equations representing the polyenergetic mode, while the red and blue lines correspond to the sequence of iterates generated by the pART and pSART algorithms, respectively. In the first experiment (left plot) the LAC values at the reference energy are $t_1 = 0.1$ and $t_2 = 0.16$, while in the left plot, $t_2 = 0.24$.

other parameters of the experiment were the same as before. It is apparent that both iterations fail to converge to the solution; pART appears to exhibit chaotic behaviour about the solution, while pSART converges to a two-cycle. Even if the iteration is started very close to the solution $\mathbf{t}^* = (0.1, 0.24)^T$, both pART and pSART diverge.

A direct computation of the 2×2 Jacobian matrix (14) for these two experiments reveals that the spectral radius of J_F , evaluated at the solution \mathbf{t}^* , is roughly 0.89 for the first case ($t_2 = 0.16$) and 1.02 for the second case ($t_2 = 0.24$). This explains why the first iteration converges, but not the second. Some further investigation reveals that this is due in large part to the discontinuities in $\frac{\partial \mu}{\partial t}$. In our experiment the reference materials were air, fat, soft tissue and bone, with tabulated LAC values at the reference energy of 70 keV equal to $0, 0.1782, 0.2033$ and 0.4948 cm^{-1} , respectively. Thus the partial derivative $\frac{\partial \mu}{\partial t}$ has a larger value for $t_2 = 0.24$ (which lies




FIG. 5. pSART iterations for $t_2 = 0.203$ (left) and $t_2 = 0.204$ (right), narrowly centred on the solution. The first set of iterations converges to solution while the second converges to a two-cycle.

between soft tissue and bone) than for $t_2 = 0.16$ (which lies between fat and air). In Fig. 5 we show the result of two more experiments where t_2 was set to values of 0.203 and 0.204 , which lie on either side of the tabulated value for soft tissue, where the derivative is discontinuous. The pSART iteration converges to the true solution in the first case but not in the second case, where it reaches a two-cycle between two points lying close to the true solution. Direct computation of the spectral radius confirms that it is equal to 0.87 in the first case and 1.16 in the second.

Fig. 6 gives a convergence map for this test case as a function of t_1 and t_2 . The effect of the discontinuities in $\frac{\partial \mu}{\partial t}$ is clearly visible in the discontinuities in $\rho(J_F)$ that occur at the values of 0.1782 – the LAC value of fat at 70 keV – and 0.2033 , the LAC value of soft tissue at 70 keV. It is apparent that the iteration transitions between convergent and non-convergent states at values of (t_1, t_2) that do not lie along these discontinuities as well. We note that the figure is not symmetric along the line $t_1 = t_2$; for example, all of the test cases that have been considered in this section would converge if the values of t_1 and t_2 were interchanged. This figure is specific to the system matrix that arises from the ray paths illustrated in Fig. 1, and would be different for other paths.

III. NUMERICAL EXPERIMENTS

We have shown in the previous section that the pSART iteration is not guaranteed to converge in general, despite the success of the algorithm demonstrated in Ref. 18. One possible explanation is that the system matrices that arise in CT imaging are typically quite sparse and structured, compared to the 2×2 matrix that was considered in the previous section. Thus, it is worth investigating whether the spectral radius of the Jacobian matrix of pSART is likely to exceed one for more realistic CT system matrices.

In the following numerical experiment we consider the problem of reconstructing an $N \times N$ pixel image for $N = 100, 200, 400$ and 800 . We simulate parallel beam data




FIG. 6. Spectral radius and convergence maps for the 2×2 polyenergetic experiment. Left figure shows the spectral radius $\rho(J_F)$ as a function of t_1 and t_2 ; right figure is the equivalent binary map indicating whether the iteration converges (blue) or not (red). White crosses indicate the cases shown in Figs. 4 and 5.




FIG. 7. FORBILD phantom for the case $N = 800$. Color window is restricted to $[0.195, 0.215]$ to show low-contrast features; bony structures (white) have a LAC of 0.4948 at the reference energy of 70 keV.

acquired at m equally spaced views over 180° , with $m = 180, 360, 720$ and 1440 , respectively. Forward projection (multiplication by A) is implemented using the `radon` command in Matlab, while backward projection (multiplication by A^T) uses `iradon` with no filtering. Since $J_F(\mathbf{t}^*)$ depends on the object \mathbf{t}^* that we wish to reconstruct, we must consider a specific object to analyze the convergence of pSART. We use an $N \times N$ slice of the FORBILD numerical head phantom [22], which consists of bone and soft tissue, and includes some low-contrast features. An image of the phantom for the case $N = 800$ is shown in Fig. 7.

With these elements in place, we can approximate the spectral radius of the matrices associated with the SART iteration (11) and the Jacobian matrix of the pSART iteration (14, 15). To approximate the spectral radius, we use the power iteration (see for instance Ref. 23), which provides iterative estimates of the largest eigenvalue of a matrix M and the associated eigenvector, denoted by $\lambda^{(k)}$ and $\mathbf{x}^{(k)}$, respectively. The iterations were started from a random unit vector $\mathbf{x}^{(0)}$ and run until the largest element of the residual, $\|M\mathbf{x}^{(k)} - \lambda^{(k)}\mathbf{x}^{(k)}\|_\infty$, was less than 10^{-4} .

Table I shows the computed approximations to $\rho(T)$

TABLE I. Results of the power iteration for the matrix T used for the SART iteration, and the Jacobian matrix of the pSART iteration, for different phantom sizes. N is the dimension the image to be reconstructed and m is the total number of views. $\rho(T)$ is the spectral radius of the matrix T , and $\rho(J_F(\mathbf{t}^*))$ is the spectral radius of the Jacobian of pSART iteration, evaluated for the $N \times N$ FORBILD phantom. The number of iterations run for each power iteration are displayed below the estimate of the spectral radius.

| N | m | $\rho(T)$ | $\rho(J_F(\mathbf{t}^*))$ |
|-----|------|-----------|---------------------------|
| 100 | 180 | 0.999957 | 0.999960 |
| | | 2535 its | 2802 its |
| 200 | 360 | 0.999957 | 0.999958 |
| | | 2899 its | 2880 its |
| 400 | 720 | 0.999952 | 0.999954 |
| | | 3475 its | 3544 its |
| 800 | 1440 | 0.999933 | 0.999935 |
| | | 3544 its | 3595 its |

and $\rho(J_F(\mathbf{t}^*))$ for the different test cases, along with the number of power iterations required to obtain the estimate. It is apparent that there is virtually no difference between the spectral radius of the SART iteration matrix, T , and that of the pSART Jacobian matrix, $J_F(\mathbf{t}^*)$, in any of the studied cases. In no instances did the computed estimate of $J_F(\mathbf{t}^*)$ ever exceed one, which would cause the iteration to diverge in the neighbourhood of the solution. We conclude that for more realistic CT imaging scenarios, the pSART iteration is, at the very least, likely to exhibit local convergence in the neighbourhood of the solution, with a comparable rate of convergence to SART. We note, however, that the spectral radius of both the SART and pSART iterations is very close to one, indicating that the convergence will be slow in the neighbourhood of the solution. The convergence can likely be accelerated by using subsets of the projection data, as was proposed in the original SART algorithm [4].

This analysis establishes only local convergence in the neighbourhood of the solution. Global convergence (i.e. from an arbitrary initial estimate) is more difficult to establish in general, and it is not obvious whether there exist conditions on the system matrix A , beam spectrum, choice of reference energy, etc. to guarantee global convergence of pSART. This is fairly typical of nonlinear fixed-point iterations, Newton's method being a well-known example. The method does appear to be fairly robust with respect to the choice of starting point, however, as the images in Ref. 18 were produced from an initial estimate consisting only of zeros; our own experience with the algorithm confirms that this choice of initial estimate works well.

Additionally, our analysis assumes the existence of an exact solution, whereas in practice the system is likely to be inconsistent due to factors such as noisy measurements and model mismatch. As with any reconstruction

algorithm, model mismatch will produce artifacts in the reconstructed image; Ref. 18 includes some experiments quantifying the effect of spectrum error. In the presence of noisy data, our experience indicates that pSART exhibits the “semi-convergence” behaviour typical of other iterative methods (see e.g. Ref. 24, p.89); namely, that the algorithm initially converges towards the solution, but the image eventually deteriorates with further iterations due to the effects of noise. As noted in Ref. 18, this problem could potentially be addressed with the use of statistical modeling or edge-preserving regularization.

CONCLUSIONS

In this paper we have analyzed the convergence of a recently proposed polyenergetic SART (pSART) algorithm. We show that the spectral radius of the Jacobian of the nonlinear pSART iteration may be larger than one in some cases. Thus the method is not mathematically guaranteed to converge to a solution of the nonlinear system of polyenergetic equations, in general. For system matrices of the type encountered in CT imaging, however, our empirical results indicate that the spectral radius of the Jacobian matrix, evaluated for a prototypical head phantom, is essentially the same as the spectral radius of the matrix corresponding to the convergent SART iteration. Thus in practice it seems that the method is likely to converge at roughly the same rate as SART.

ACKNOWLEDGMENTS

The author thanks Adel Faridani (Oregon State University) and Yuan Lin (Duke University) for helpful discussion on this paper, and Federico Poloni (Università di Pisa) for helpful suggestions on the proof of Lemma A.2 in the Appendix.

APPENDIX

We provide a short proof of convergence for SART subject to a condition on the rank of A . Convergence of SART has been established previously in Refs. 6 and 7. This proof is somewhat shorter and is intended to complement the analysis we have presented for pSART.

Let $T = I - DA^TMA$, where A is the $m \times n$ system matrix and D and M are defined in (4). We assume that A has rank n , meaning that there are as many linearly independent equations as there are unknowns. This implies that $m \geq n$, i.e. that there are at least as many measurements as there are unknowns. The SART iteration then has the form

$$\mathbf{x}^{(k+1)} = T\mathbf{x}^{(k)} + \mathbf{c}$$

We first prove two lemmas.

Lemma A.1: *Let $W = DA^TMA$. Then, $\rho(W) \leq 1$, with equality if all elements of A are positive.*

Proof: A direct calculation gives the (i, j) th element of W as

$$w_{ij} = \frac{1}{\beta_i} \sum_{k=1}^m \frac{a_{ki}a_{kj}}{\gamma_k}.$$

It follows that the sum of the absolute values in row i of W is:

$$\begin{aligned} \sum_{j=1}^n |w_{ij}| &= \sum_{j=1}^n \frac{1}{\beta_i} \left| \sum_{k=1}^m \frac{a_{ki}a_{kj}}{\gamma_k} \right| \\ &\leq \frac{1}{\beta_i} \sum_{k=1}^m \frac{|a_{ki}|}{\gamma_k} \sum_{j=1}^n |a_{kj}| \\ &= \frac{1}{\beta_i} \sum_{k=1}^m \frac{|a_{ki}|}{\gamma_k} \gamma_k \\ &= \frac{1}{\beta_i} \sum_{k=1}^m |a_{ki}| \\ &= \frac{1}{\beta_i} \beta_i \\ &= 1 \end{aligned}$$

Since the max norm of a matrix, $\|\cdot\|_\infty$, is equal to the maximum row sum, and the spectral radius of a matrix cannot be greater than the max norm, it must be true that

$$\rho(W) \leq \|W\|_\infty \leq 1.$$

When all elements of A are positive (which is the case in tomographic applications), every row of W sums exactly to 1, and so $\lambda = 1$ is an eigenvalue of W (with the associated eigenvector consisting of all ones), and $\rho(W) = 1$.

On its own, this lemma only tells us that the eigenvalues of W have magnitude less than 1. We also need the following result:

Lemma A.2: *All eigenvalues of W are positive real numbers.*

Proof: This Lemma is an application of Theorem 7.6.3 from Ref. 25. We first remind the reader of two definitions:

1. Two $n \times n$ matrices A and B are *similar* if there exists an invertible matrix P such that $A = PBP^{-1}$. If A and B are similar, then they have exactly the same eigenvalues.
2. Two real $n \times n$ matrices A and B are *congruent* if there exists an invertible matrix P such that $A = PBP^T$. If A and B are symmetric, then Sylvester’s law of inertia states that they have the same inertia, meaning the same number of positive, negative, and zero eigenvalues. (Recall that the eigenvalues of a symmetric matrix are always real numbers).

Now, let $V = A^T M A$. Then, V is a symmetric positive semidefinite $n \times n$ matrix, since it can be written as the product of a matrix and its transpose: $V = (M^{\frac{1}{2}} A)^T (M^{\frac{1}{2}} A)$. Furthermore, since A has rank n and $M^{\frac{1}{2}}$ is a diagonal matrix, $M^{\frac{1}{2}} A$ has rank n and so does V . It follows that V is invertible, and so zero is not an eigenvalue of V . Thus all eigenvalues of V are positive. We then have the following two results:

1. $W = DV$ is similar to $D^{\frac{1}{2}} V D^{\frac{1}{2}}$, since $D^{\frac{1}{2}} V D^{\frac{1}{2}} = D^{-\frac{1}{2}} D V D^{\frac{1}{2}}$.
2. $D^{\frac{1}{2}} V D^{\frac{1}{2}}$ is congruent to V .

The first result implies that the eigenvalues of W are the same as the eigenvalues of $D^{\frac{1}{2}} V D^{\frac{1}{2}}$, while the second implies that the eigenvalues of this matrix must have the same signs as the eigenvalues of V . So, since all eigenvalues of V are positive, all eigenvalues of W must be positive as well.

Theorem A.3: The matrix T satisfies $\rho(T) < 1$, and hence the SART iteration converges.

Proof: Lemmas A.1 and A.2 prove that any eigenvalues λ of W satisfy $0 < \lambda \leq 1$. Thus the eigenvalues of $T = I - W$ satisfy $0 \leq \lambda < 1$, and so $\rho(T) < 1$.

-
- [1] S. Kaczmarz. Approximate solution of systems of linear equations (trans. P.C. Parks). *International Journal of Control*, 57(3):1269–1271, 1993. (Originally published as: Angenäherte Auflösung von Systemen linearer Gleichungen. Bulletin International de l'Academie Polonaise des Sciences. Lett A, 355-357. 1937).
 - [2] R. Gordon, R. Bender, and G. T. Herman. Algebraic reconstruction techniques (ART) for three-dimensional electron microscopy and X-ray photography. *Journal of theoretical biology*, 29(3):471–481, 1970.
 - [3] K. Tanabe. Projection method for solving a singular system of linear equations and its applications. *Numerische Mathematik*, 17(3):203–214, 1971.
 - [4] A. H. Andersen and A. C. Kak. Simultaneous algebraic reconstruction technique (SART): a superior implementation of the ART algorithm. *Ultrasonic Imaging*, 6(1):81–94, 1984.
 - [5] A.C. Kak and M. Slaney. *Principles of Computerized Tomographic Imaging*, chapter 7. SIAM, 2001.
 - [6] Y. Censor and T. Elfving. Block-iterative algorithms with diagonally scaled oblique projections for the linear feasibility problem. *SIAM Journal on Matrix Analysis and Applications*, 24(1):40–58, 2002.
 - [7] M. Jiang and G. Wang. Convergence of the simultaneous algebraic reconstruction technique (SART). *IEEE Transactions on Image Processing*, 12(8):957–961, 2003.
 - [8] R. A. Brooks and G. Di Chiro. Beam hardening in X-ray reconstructive tomography. *Phys. Med. Biol.*, 21(3):390, 1976.
 - [9] J. Hsieh, R. Molthen, C. Dawson, and R. Johnson. An iterative approach to the beam hardening correction in cone beam CT. *Med. Phys.*, 27(1):23–29, 2000.
 - [10] C. H. Yan, R. T. Whalen, G. S. Beaupré, S. Y. Yen, and S. Napel. Reconstruction algorithm for polychromatic CT imaging: application to beam hardening correction. *IEEE Trans. Med. Imag.*, 19(1), 2000.
 - [11] B. De Man, J. Nuyts, P. Dupont, G. Marchal, and P. Suetens. An iterative maximum-likelihood polychromatic algorithm for CT. *IEEE Trans. Med. Imag.*, 20(10):999–1008, 2001.
 - [12] I. A. Elbakri and J. A. Fessler. Statistical image reconstruction for polyenergetic X-ray computed tomography. *IEEE Trans. Med. Imag.*, 21(2):89–99, 2002.
 - [13] I. A. Elbakri and J. A. Fessler. Segmentation-free statistical image reconstruction for polyenergetic X-ray computed tomography with experimental validation. *Phys. Med. Biol.*, 48:2453–2477–99, 2003.
 - [14] G. Van Gompel, K. Van Slambrouck, M. Defrise, K.J. Batenburg, J. de Mey, J. Sijbers, and J. Nuyts. Iterative correction of beam hardening artifacts in CT. *Med. Phys.*, 38(7):S36–S49, 2011.
 - [15] N. Rezvani. *Iterative Reconstruction Algorithms for Polyenergetic X-Ray Computerized Tomography*. PhD thesis, University of Toronto, 2012.
 - [16] T. Humphries and A. Faridani. Segmentation-free quasi-newton method for polyenergetic CT reconstruction. In *2014 IEEE Nuclear Science Symposium Conference Record*, 2014.
 - [17] Y. Lin and E. Samei. A fast poly-energetic iterative FBP algorithm. *Phys. Med. Biol.*, 59:1655–1678, 2014.
 - [18] Y. Lin and E. Samei. An efficient polyenergetic SART (pSART) reconstruction algorithm for quantitative myocardial CT perfusion. *Med. Phys.*, 41(2):021911–1 – 021911–14, 2014.
 - [19] Spektrum – Siemens OEM Products. <https://w9.siemens.com/cms/oemproducts/home/x-raytoolbox/spektrum/pages/default.aspx>.
 - [20] J. M. Boone and J. A. Seibert. An accurate method for computer-generating tungsten anode x-ray spectra from 30 to 140 kV. *Med. Phys.*, 24(11):1661–1670, 1997.
 - [21] J. H. Hubbell and S. M. Seltzer. Tables of X-ray mass attenuation coefficients and mass energy-absorption coefficients. *National Institute of Standards and Technology*, 1996.
 - [22] Z. Yu, F. Noo, F. Dennerlein, A. Wunderlich, G. Lauritsch, and J. Hornegger. Simulation tools for two-dimensional experiments in X-ray computed tomography using the FORBILD head phantom. *Phys. Med. Biol.*, 57(13):N237–N252, 2012.
 - [23] L. N. Trefethen and D. Bau. *Numerical Linear Algebra*, Chapter 27: Rayleigh Quotient, Inverse Iteration. SIAM, 1997.
 - [24] F. Natterer. *The Mathematics of Computerized Tomography*. Springer, 1986.
 - [25] R. A. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge University Press, 1985.